




VISIUM

[A PRACTICAL ROADMAP]

The C-level guide to kickstart your AI journey


Content overview

Information for all board members	{ 3 }
AI glossary	{ 5 }
Implications for the CIO function	{ 9 }
Potential AI use cases	{ 11 }
Case studies from Visium	{ 12 }

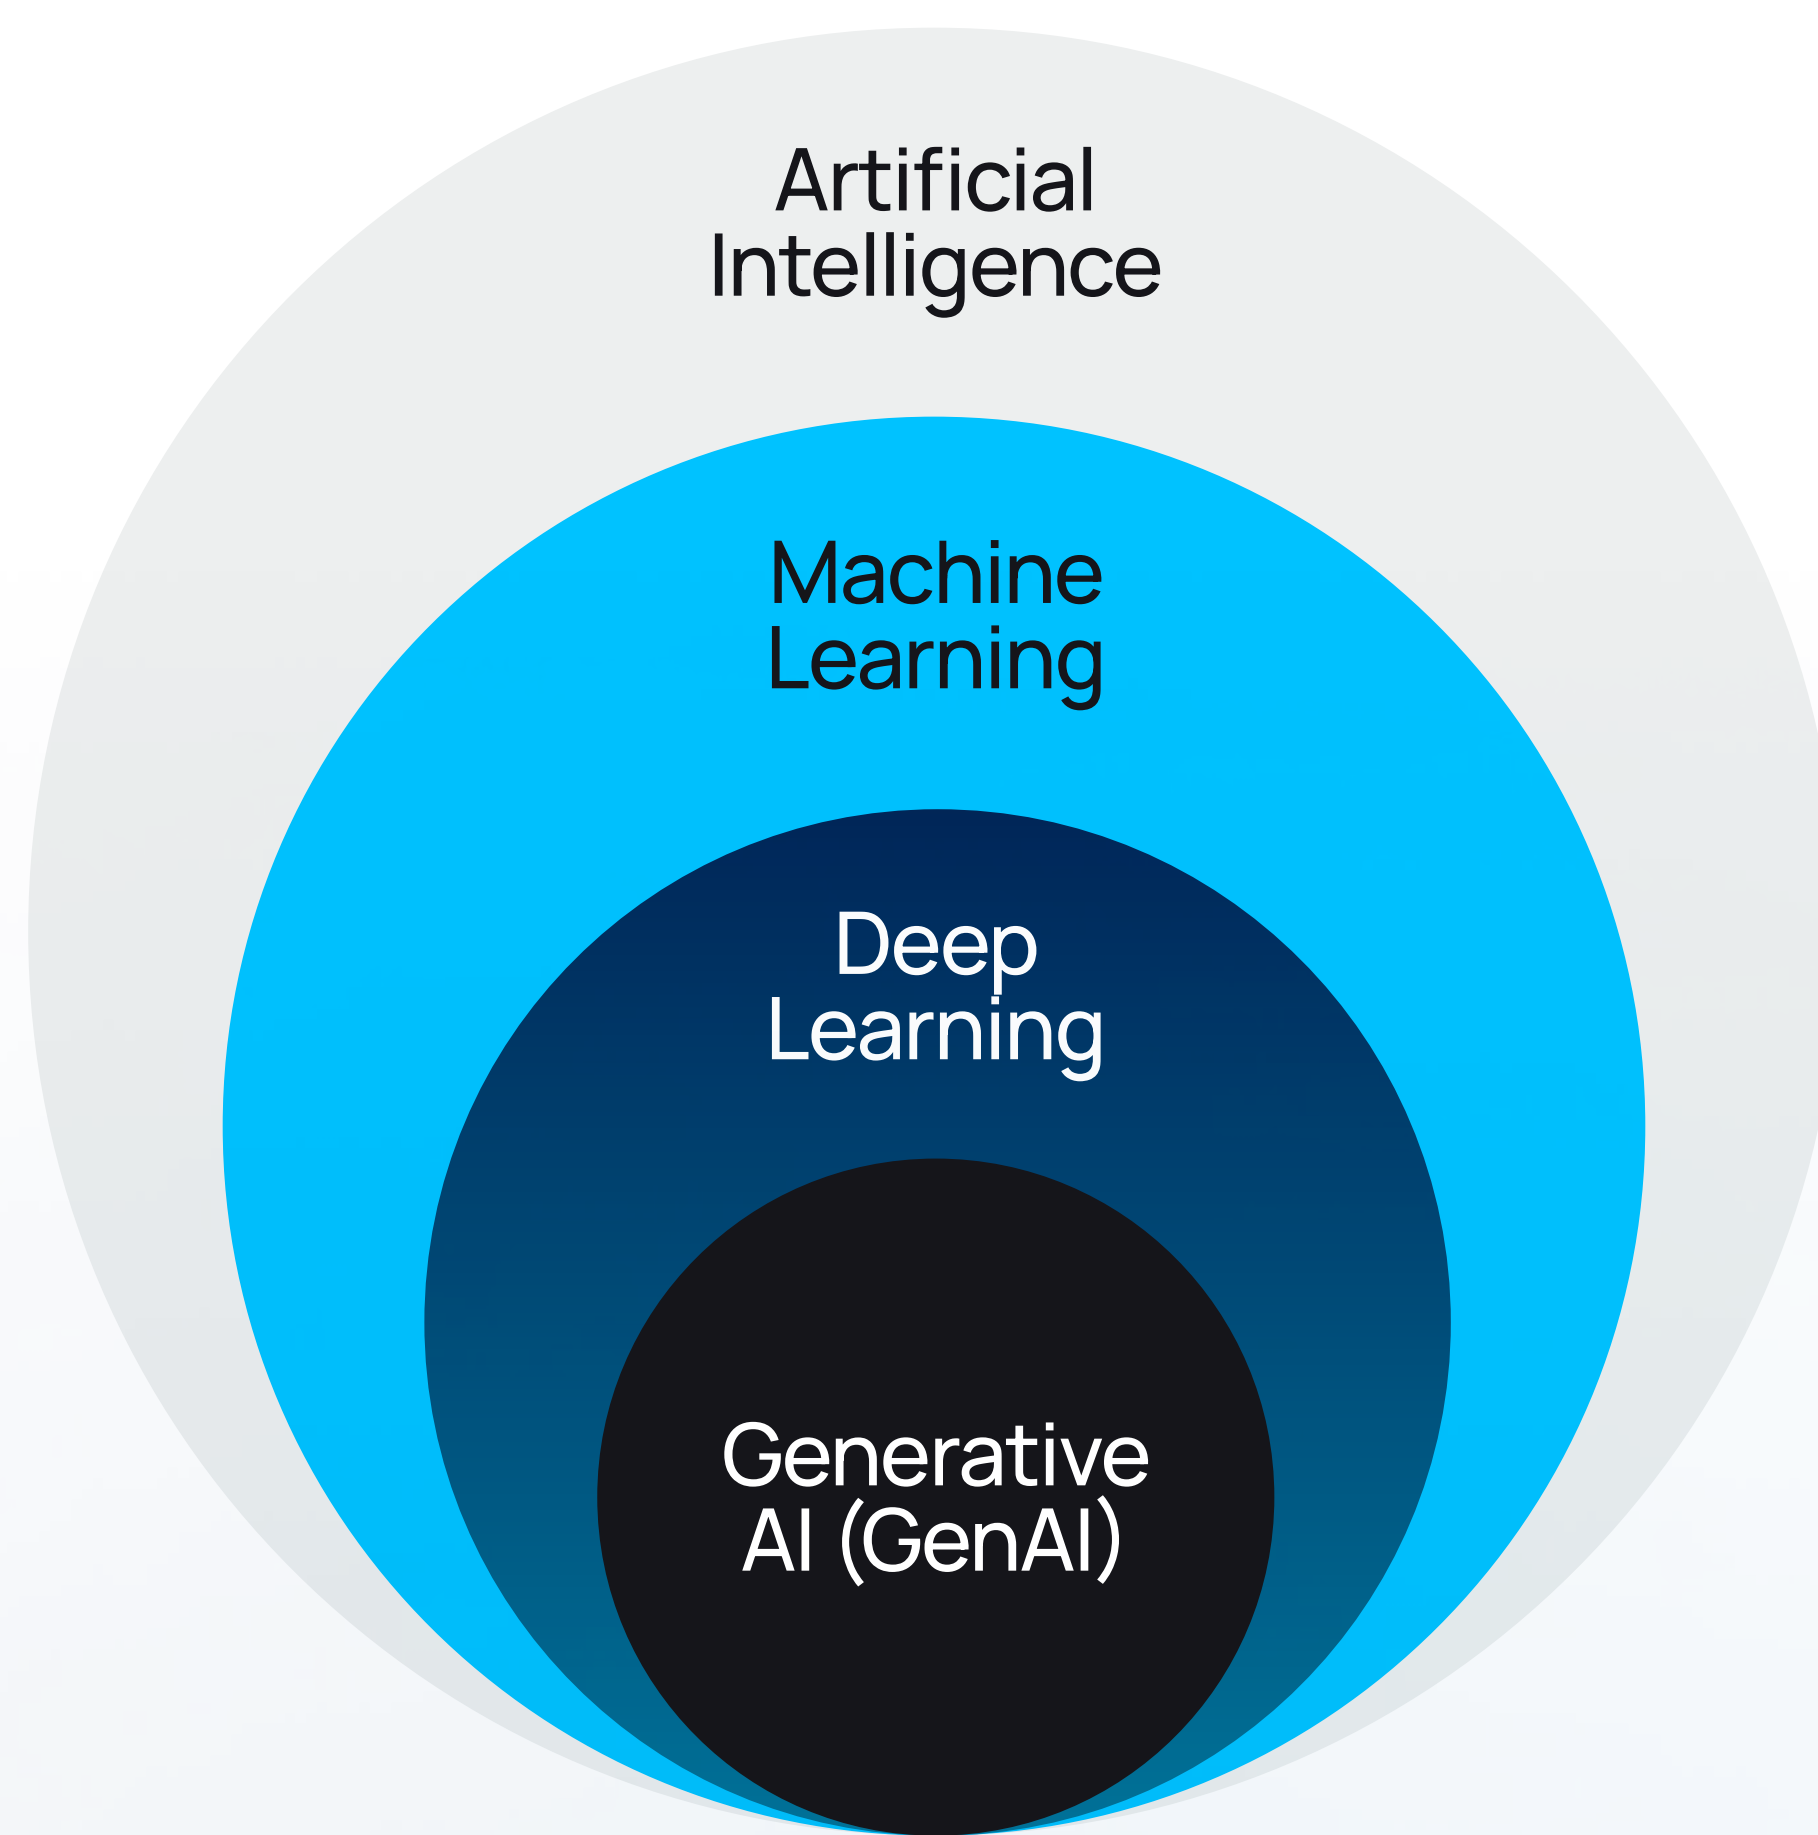


Many companies feel overwhelmed when beginning their AI journey, and as a leader, you may feel the same amid a sea of information and recommendations.

This guide will help you cut through the noise and navigate your AI journey one step at a time.



AI glossary



Artificial Intelligence (AI)

Refers to a field of computer science that creates machines capable of human-like tasks, such as perception, reasoning, planning, learning, and decision-making.


Example: AI optimizes scheduling for tasks like employee shifts or transportation.



Machine Learning (ML)

A branch of AI where machines learn to solve problems by generalizing patterns from provided examples.

Example: ML detects patterns and anomalies to identify fraud.



Deep Learning

A specialized area within machine learning using brain-like neural networks to process complex data such as text or images.

Example: Neural networks process camera sensory data for autonomous driving.



Generative AI (GenAI)

A category of deep learning algorithms focused on creating realistic content from data such as text, images, or audio.

Example: Popular examples are ChatGPT (for text) and DALL-E (for images).

AI Glossary

Understand AI - from data to decisions

Unlike programming that operates on deterministic rules and predefined instructions, AI, particularly Machine Learning (ML), uses a probabilistic approach. Instead of following rigid rules, AI systems learn patterns from data, make predictions, and adapt over time based on feedback and observed outcomes. This adaptive, data-driven method allows AI to handle uncertainty and continuously improve, making it fundamentally different from the static, rule-bound nature of traditional computing.

AI's data dependency - balancing data and expert knowledge

The data needs of AI projects vary based on the specific problem and algorithm used. While having large amounts of labeled data can reduce the need for some types of expertise, expert knowledge remains key in building relevant and usable AI solutions. Experts provide the domain-specific insights needed to structure data, define meaningful objectives, and interpret results accurately. When data is limited, expert involvement becomes even more critical to ensure that AI models are not just technically sound but also tailored to the needs of end users.

Harnessing AI - key metrics for model performance and business value

The true measure of an AI model's success is measured by its technical performance but also by the business value it creates. To effectively assess AI models, it's important to distinguish between model-specific metrics and those that reflect broader business impact. For example, in a classification task like defect detection in quality control, model-specific metrics such as false positives (detecting a defect when none exists) and false negatives (missing an actual defect) are crucial for understanding model performance. However, the ultimate goal is to connect these technical metrics with business-relevant outcomes like operational efficiency and key financial outcomes like cost reduction, revenue growth, and improved profitability.

AI Glossary

Humans AI - augmentation, not replacement

While AI can mimic specific human tasks, it is still just one tool in the broader toolbox for creating value. AI can assist in thinking, planning, and decision-making processes, but it relies on human guidance to define, refine, and apply these capabilities effectively. This distinction is crucial for leaders to understand; AI accelerates and amplifies what humans can achieve but does not replace the nuanced judgment, creativity, and contextual understanding that only humans bring.

Scaling AI - the need for full pipeline support

Unlike traditional software, AI cannot be separated from its data, making it hard to break the system into smaller modules that can be easily managed or updated separately. This challenge is amplified by the lack of specialized tools for managing AI workflows, often turning simple AI pilots into unexpectedly complex scaling efforts, sometimes costing 10-20 times more than the pilot phase. Scaling AI requires continuous management of the entire so-called ML pipeline — a process that oversees every step from data preparation and model training to deployment and user interaction. Mastering AI at scale remains one of the biggest hurdles in business today.


Ownership of AI - make vs. buy

The decision to develop AI in-house or purchase an external solution isn't straightforward due to factors like strategic value, data ownership, customization, performance of external solutions, and total cost of ownership. In-house development offers control and tailored solutions but requires substantial investment in expertise, infrastructure, and continuous maintenance. External solutions provide quicker access to advanced capabilities and ongoing vendor support but may involve compromises on data privacy, integration complexities, and adaptability to specific needs. These implications, especially with their strong financial component, must be carefully considered to align with your business strategy.

Navigating AI ethics - mitigation bias

Ethical AI goes beyond compliance; it's about building systems that are fair, transparent, and aligned with your company's values. Addressing bias is crucial, as AI decisions are deeply tied to the quality and biases of its training data. Principles like 'equal treatment' require human intervention to define fairness, making us rethink whether treating groups 'on average' is sufficient. This underscores the critical role of human judgment as a safeguard, but also as an active guide in ensuring AI decisions are contextually aware and aligned with the broader ethical goals of the organization.

The AI opportunity



Most companies are still experimenting with AI, facing challenges in strategy and execution.

However, its impact is clear. AI can boost EBITDA by 20-30%, according to McKinsey and Accenture.

AI isn't just a technical upgrade; it transforms businesses at every level. Tools like LLM chatbots (e.g., ChatGPT) are changing how we work and learn. For companies, AI drives efficiency, sparks innovation, and reshapes operations, enhancing competitiveness. At the industry level, AI optimizes processes from R&D to distribution, redefining whole value.

To keep in mind

AI projects differ from traditional IT initiatives due to their experimental nature. The outcomes are not always guaranteed, and success often hinges on iterative experimentation and adaptation. Effective AI adoption requires significant data science expertise, strategic planning, and a shift in how leaders engage with data science.

Implications for the CIO function



Implications for the CIO function

Core responsibilities

01 Data governance and infrastructure

As the architect of the company's AI foundation, the CIO shapes the data landscape that drives AI success. This begins with a robust data strategy that prioritizes data acquisition, integration, and accessibility (e.g., ensuring your data is findable, accessible, cross-functional, and reusable). Ensuring the availability of high-quality labeled data, particularly for supervised machine learning, is essential. The CIO must develop efficient data pipelines and scalable AI architectures, working closely with business units across the organization to align technology with strategic goals.

02 Risk management

The CIO plays a critical role in evaluating risks facing the company. While AI is highly effective at identifying various risks, it also brings its own set of challenges, including potential liabilities, regulatory compliance concerns, and legal complexities. While AI actions are measurable and can often be adapted to fit regulatory frameworks better than human actions, inspecting and explaining each decision can be challenging without prior setup. For the CIO, it's essential to ensure compliance through proactive measures, such as setting up predefined monitoring protocols and engaging in preemptive discussions with regulators.

03 Applications at scale

Deploying AI at scale (i.e. MLOps and LLMOps) is key to unlocking its full potential and driving substantial value for the company. Beyond ensuring robust systems and infrastructure, the CIO plays a strategic role in aligning technology with business needs, facilitating cross-functional collaboration, and integrating AI solutions seamlessly into operations. Successful scaling requires the CIO to lead the charge in bridging technical capabilities with organizational readiness, ensuring that AI initiatives are executed efficiently and effectively across the enterprise.

04 Make-or-buy

Lastly, the CIO plays a key role in deciding whether to develop AI solutions in-house or collaborate with external partners. Unlike traditional IT projects, AI initiatives often demand new and highly tailored collaboration models. Typically, the CIO has the most expertise in AI among the management team and serves as a valuable advisor for the rest of the board. Integrating AI with company data brings added complexities in IP ownership, contract management, and sourcing, which requires the CIO to closely engage with AI partners and be prepared to handle these unique demands.

Implications for the CIO function

• Example of AI in action: Use cases

Use cases	Description
Centralized data platform	Streamline data access, automatically catalog and tag data from many sources. A centralized data platform breaks down silos, enhances governance, ensuring your teams have timely, secure access to the data they need.
Data quality monitoring	AI monitors data quality in real-time, detecting anomalies, duplicates, and errors to ensure reliable data for AI models, helping the CIO uphold data standards.
AI-powered knowledge center	An AI-driven RAG system can centralize past learnings, documentation, and insights. For a CIO, this means turning your organization's historical data into actionable knowledge, reducing search time, enhancing decision-making, and enabling IT teams to solve problems faster with context-aware answers.

Implementing data-centric systems

🎯 The challenge

A leading luxury company faced fragmented, outdated data systems without central governance. Departments worked in silos, each handling data transformations differently, leading to discrepancies in insights and inconsistent decision-making due to varying terminology and calculation methods.

⚙️ The solution

Our solution involved implementing a modern Data Platform and DataOps processes to streamline data ingestion, transformation, and consumption. A comprehensive Data Governance initiative was introduced to ensure data quality, consistency, and security, across the organization.

✅ The outcome

- Two MVPs successfully deployed, meeting defined success criteria.
- Users embraced the new platform, driving the creation of additional use cases.
- Data transformations became 5x faster, significantly enhancing operational speed and accuracy.

Deploying AI with MLOps integration

🎯 The challenge

Our client built their first successful POC with Machine Learning models but was missing a comprehensive strategy and frameworks to put the models into production. They reached out for help in defining the right processes and technologies to break down data silos, allowing them to use the ML models across their multiple sites and drive business value.

🚀 The solution

We assessed our client's business and technical needs and proposed a detailed tech stack along with an MLOps process to ensure effective Model Lifecycle Management. In addition to providing implementation examples, we assisted with setting up and migrating their internal stack, all while coaching their teams on using the new tools.


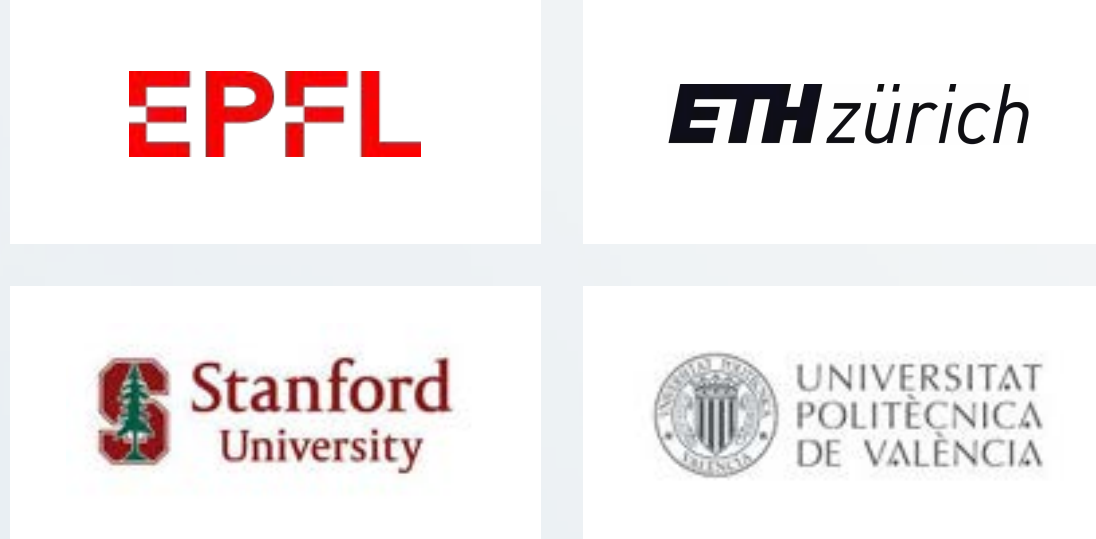
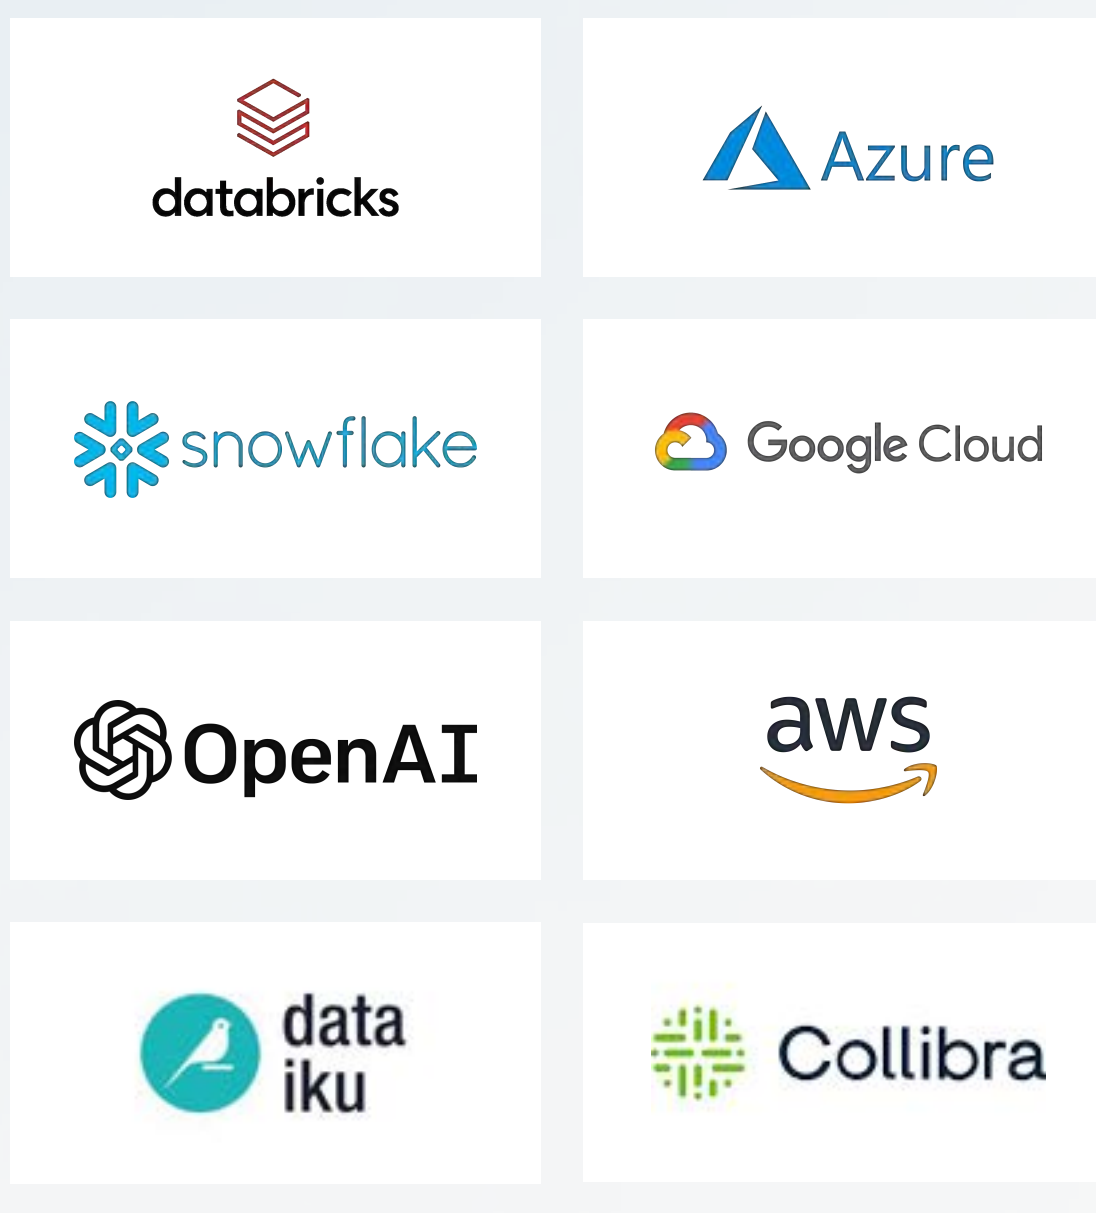



✅ The outcome

- Increased operationalized ML models from 0 to 10.
- Unlocked access to data and resources across multiple locations, improving collaboration.
- Established a strong foundation for executing future strategic AI projects.

About Visium

Visium is a leading Swiss AI & Data consultancy helping enterprises turn their most strategic AI initiatives into measurable business outcomes. Since 2018, we've delivered over 250 solutions for global leaders like Roche, Novartis, Nestlé, and dsm-Firmenich; helping them drive efficiency, unlock new revenue streams, and scale solutions that deliver long-term value.

Built on a foundation of ethical innovation, our mission is to make AI work for business, people, and long-term impact.

70+ Engineers and consultants throughout Europe	50+ Happy enterprise clients 	Our partners Academia  Technology 
250+ AI & data engagements		
Pan-European presence 	Recognized leaders 2nd fastest growing company in Europe by Financial Times 	

Get in touch

✉ contact@visium.com

🌐 www.visium.com



AI initiatives delivering real impact at scale

www.visium.com

